arXiv:1701.08075v1 [quant-ph] 25 Jan 2017

Categorical Probabilistic Theories

Stefano Gogioso
University of Oxford
stefano.gogioso@cs.ox.ac.uk

Carlo Maria Scandolo
University of Oxford
carlomaria.scandolo@cs.ox.ac.uk

We present a simple categorical framework for the treatment of probabilistic theories, with the aim of reconciling the fields of Categorical Quantum Mechanics (CQM) and Operational Probabilistic Theories (OPTs). In recent years, both CQM and OPTs have found successful application to a number of areas in quantum foundations and information theory: they present many similarities, both in spirit and in formalism, but they remain separated by a number of subtle yet important differences. We attempt to bridge this gap, by adopting a minimal number of operationally motivated axioms which provide clean categorical foundations, in the style of CQM, for the treatment of the problems that OPTs are concerned with.

1 Introduction

Categorical methods are finding more and more applications in the foundations of quantum theory and quantum information, with two main frameworks currently dominating the scene: Categorical Quantum Mechanics (CQM) and Operational Probabilistic Theories (OPTs). The CQM framework [4, 7, 26, 30, 31] concerns itself with the compositional and algebraic structure of quantum system and processes, and has a heavy focus on diagrammatic methods [24, 48, 49]. The OPT framework [17, 18, 19, 21, 43, 44, 45] was developed with the aim of obtaining a characterisation of quantum theory in terms of information-theoretic axioms, inside the wider community of generalised probabilistic theories [9, 10, 11, 36, 46, 50].

Both frameworks are based on symmetric monoidal categories, and at first sight it looks like they should be complementing each other. However, a number of subtle differences have resulted in a significant disconnect between the two communities, with duplication of efforts and loss of synergy. The OPT framework is founded on explicit probabilistic structure, axioms of concrete physical inspiration and traditional proof methods, which make it hard to work with in a categorical and diagrammatic setting. Conversely, the CQM framework has clean categorical foundations and diagrammatic proof methods, but at the expense of direct physical interpretation.

We set off to reconcile the two frameworks, by proposing categorical foundations in the style of CQM for the kind of systems and results that are of interest for the OPT community. We define a probabilistic theory as a symmetric monoidal category satisfying three additional requirements: the existence of classical systems, the existence of probabilistic structure, and the possibility of defining local states. Our requirements are formulated in standard categorical terms, and can be readily adopted within the CQM framework; at the same time, they correspond to relatable physical and operational requirements that any probabilistic theory should possess. The end result is a framework which is simple and rigorous in its foundations, shows direct operational significance, and comes with more expressive power than either of the CQM or OPT frameworks alone.

It should be noted that similar efforts to connect the CQM and OPT communities are being undertaken by Ref. [53] (and ongoing work by its author): the work therein is however closer to the logical take of effectus theory, while our work has a more physical take, with classical systems at its very heart.

Submitted to:
QPL 2017

© S. Gogioso & C. M. Scandolo
This work is licensed under the Creative Commons
Attribution-Noncommercial-Share Alike License.

2 Classical Theory

The main ingredient distinguishing OPTs from the categories used in the context of CQM is the explicit presence of probabilistic mixtures, or equivalently the existence of classical systems taking part in the processes. From the point of view of OPTs, these classical systems are absorbed into the formalism, using indices and summations. From the point of view of CQM, on the other hand, the only distinction between classical and quantum systems is operational, and both should be modelled explicitly by a theory.

2.1 Classical theory as a SMC

Classical systems[1] and stochastic processes between them form a symmetric monoidal category (SMC), usually known as Stoch. The objects of Stoch are labelled by finite sets, and the morphisms $X \to Y$ in Stoch are the Y-by-X stochastic matrices, with matrix composition as sequential composition and the Kronecker product as parallel composition (aka tensor product).

From a categorical standpoint, the restriction to convex combinations required by Stoch is inconvenient: if we wish to use a summation operation, as is usually done in the treatment of classical systems and mixed-state quantum theory, we have to introduce external restrictions on the kind of sums that we allow. Alternatively, we could see the summation as defined on a larger category, with Stoch arising as the sub-category of suitably *normalised* processes. This latter approach is in line with the general philosophy of CQM: one studies a broader collection of abstract processes, easier to describe and to manipulate, and then assembles/restricts them appropriately to obtain the subcategory of concrete, physically relevant processes. As our model of classical systems and (potentially non-normalised) non-deterministic processes between them, we take the SMC $\mathbb{R}^+$-Mat: just like Stoch, the objects are the finite non-empty sets, but the morphisms $X \to Y$ are ALL the $\mathbb{R}^+$-valued matrices in $\left(\mathbb{R}^+\right)^{Y\times X}$, not just the stochastic ones. We refer to these maps as the **probabilistic processes**, and use **normalised** to denote the stochastic ones.

2.2 Linear structure

We can equip $\mathbb{R}^+$-Mat with a well-defined notion of summation by suitably enriching it in commutative monoids: we endow the morphisms $X \to Y$ with the linear structure of $\left(\mathbb{R}^+\right)^{Y\times X}$, i.e. we take matrix addition as our summation operation and the zero matrix as its neutral element (modelling the impossible process). Traditionally, the only requirement for a category enriched in commutative monoid is that sequential composition be linear[2], but in SMCs it makes sense to further require that parallel composition be linear as well[3]. This is indeed the case for our choice of enrichment.

2.3 Deterministic processes

A special place amongst the many probabilistic processes between two classical systems X and Y is held by the **deterministic processes**[4], the functions $X \to Y$. The deterministic processes form the sub-SMC fSet of $\mathbb{R}^+$-Mat, and the Kronecker product restricts to the Cartesian product on them. The states in

[1] This work, like most works in OPTs and CQM, is concerned with *finite* classical systems alone.

[2] Specifically, we require $f\circ(g+h) = (f\circ g)+(f\circ h)$, $(g+h)\circ f = (g\circ f)+(h\circ f)$, $0\circ f = 0$ and $f\circ 0 = 0$.

[3] Things are a bit more complicated with the tensor structure: further to the tensor product being linear, just like composition, but we need to require the associator and unitors to be linear as well.

[4] In line with its common meaning in computer science, we use the word **deterministic** to denote classical processes which map each definite input to a single definite output. We use **stochastic** (or normalised probabilistic, in this work) to denote classical processes which more generally map definite inputs to probability distributions over definite outputs.

fSet for a classical system X are its **deterministic states**, the probability distributions δ_x concentrated entirely on a single point $x \in X$ (or, if you prefer, the linear extensions of the points of X themselves). The singleton set $\mathbb{1} := \{1\}$ is the terminal object in fSet, and we refer to the unique effect $X \to \mathbb{1}$ as the **discarding map**[5] on the classical system X, which is graphically denoted by a ground symbol $\dashv\!\Vdash_X$. The deterministic processes sit inside a larger sub-SMC fPFun of **partial deterministic processes**, with partial functions as morphisms.

2.4 Resolution of the identity

In order to clearly distinguish classical systems from more general systems in the theory, we will adopt a graphical convention in which dashed wires are used to denote systems which are guaranteed to be classical, so that the identity morphism $X \to X$ for a classical system X will be denoted as follows:

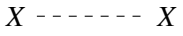

(2.1)

Using the enrichment, we obtain the following resolution of the identity:

(2.2)

The state labelled by $x \in X$ above is the deterministic state corresponding to x; the effect labelled by x above is the one sending x to $1 \in \mathbb{R}^+$ and all $x' \neq x$ to $0 \in \mathbb{R}^+$ (or, if you prefer, the linear extension of the partial function $X \rightharpoonup \mathbb{1}$ which sends x to $1 \in \mathbb{1}$ and is undefined on all other $x' \in X$).

2.5 Marginalisation

Because the resolution of the identity is in terms of deterministic states, we immediately obtain the following resolution of the discarding map:

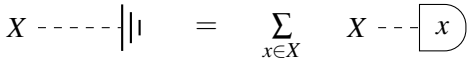

(2.3)

This means that discarding a subsystem in $\mathbb{R}^+$-Mat corresponds to the usual notion of **marginalisation**. From the CQM perspective, the process of localisation of states and processes in arbitrary SMCs is captured by the notion of **environment structure** [25, 32]. Because classical states and processes are localised by marginalising, it is no surprise that discarding maps in $\mathbb{R}^+$-Mat satisfy the requirements to form an environment structure:

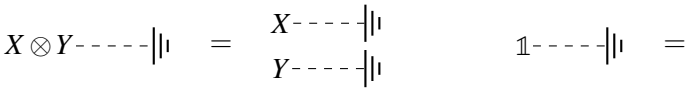

(2.4)

The empty diagram on the right hand side of the second equation is the scalar 1, which as a diagrammatic convention is usually omitted.

[5] In a not-at-all-unexpected coincidence, it makes total sense to call this **the deterministic effect**.

2.6 Normalised processes

A choice of environment structure always singles out a sub-SMC of **normalised** processes, namely those processes f satisfying the following condition:

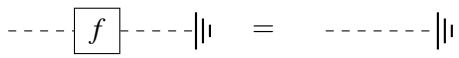

(2.5)

The abstract normalisation condition above allows us to recover Stoch in a categorically fashionable way as the sub-SMC of normalised processes in $\mathbb{R}^+$-Mat. As a bonus, we also get fSet as the sub-SMC of normalised processes in fPFun.

2.7 Probabilities as a resource

In computer science, it is common for commutative semirings to model resources used in computation: for example, the *probability semiring* $(\mathbb{R}^+, +, \times)$ is used to model probabilistic computation, the *boolean semiring* $(\mathbb{B}, \vee, \wedge)$ is used to model existence problems, and the *natural numbers* $(\mathbb{N}, +, \times)$ are used to model counting problems[6]. A similar approach could be adopted in physics: probabilities (corresponding to the semiring $\mathbb{R}^+$) could be seen as a resource providing non-determinism in classical systems, and we might wish to study the toy theories obtained by using other resources (corresponding to other semirings) in their place. This is, for example, the road taken by the sheaf-theoretic framework of Ref. [2], where a number of important results on non-locality and contextuality are obtained by confronting the probabilities (the semiring $\mathbb{R}^+$) with possibilities [5] (the semiring $\mathbb{B}$) and signed probabilities [6, 39] (the semiring $\mathbb{R}$, which is given an operational interpretation in Ref. [3]).

If we use a generic commutative semiring R to model non-determinism, the SMC of classical systems and **R-probabilistic** processes between them is given by the category R-Mat of R-valued matrices: we refer to the normalised states as **R-distributions**, and to the normalised processes as **R-stochastic**. As an example, $\mathbb{B}$-Mat is the category fRel of finite sets and relations between them, a well-studied toy model in CQM [27, 33, 40]. All the theory we have developed above for $\mathbb{R}^+$-Mat straightforwardly extends to the general semirings, as long as we replace Stoch with an appropriate notion of normalised processes. The categories fSet and fPFun of total and partial deterministic processes are always sub-SMCs of R-Mat, showing that our interpretation of R as modelling a notion of non-determinism is sound.

3 Probabilistic Theories

When talking about a *theory*, we broadly mean a categorical model (a SMC) that captures physical systems and the compositional structure of processes between them. In the spirit of CQM, we avoid restricting our attention to physical processes alone, but instead allow the presence of a whole spectrum of idealised, abstract processes that provide building blocks for physical processes, or otherwise help in reasoning about them. Contrary to the OPT framework, we take the view that a theory should specify objects and processes in a purely compositional way, and need not necessarily make immediate reference to probabilistic outcomes of tests and measurement: as a consequence, we do not take any quotient with respect to probabilistic outcomes, and we allow for the possibility that a theory specifies processes which are different but cannot be distinguished by means of tests or measurements alone (this means that some different "hidden variable theories" with the same operational predictions can be modelled in this framework).

[6] One could also mention the *tropical semiring* $(\mathbb{R}, \min, +)$ used in the Floyd-Warshall algorithm, or the *Viterbi semiring* $([0,1], \max, \times)$ used in the Viterbi algorithm.

3.1 Probabilistic theories

When talking about a *probabilistic theory*, we mean a SMC which includes at least the classical systems amongst its ranks, and which is furthermore compatible with a couple of basic operational features, namely their probabilistic structure and marginalisation. We briefly motivate our requirements below.

1. Every probabilistic theory has classical probabilistic systems under the hood: because these are themselves physical systems, we model them explicitly. In particular, their interface with other systems (e.g. measurements and preparations) can be talked about in compositional terms.
2. It makes no sense to talk about a probabilistic theory if the probabilistic structure does not extend from classical systems to arbitrary systems. If this were not the case, one would not necessarily be able to work with scenarios in which multiplexed processes are controlled by a classical random variable, or to condition a process based on a classical output.
3. Sensible theories without a notion of local state do exist (e.g. Everettian quantum theory), but they present a number of operational challenges, and they are not probabilistic in nature. Indeed, the absence of a notion of local state compatible with marginalisation makes most multi-party measurement scenarios and protocol specifications meaningless, a fate they share with the notion of no-signalling and with the probabilistic foundations of thermodynamics.

We now proceed to formalise these requirements in categorical terms.

Definition 1 (**Probabilistic Theory**)**.**
*A **probabilistic theory** is a symmetric monoidal category $\mathscr{C}$ which satisfies the following requirements.*

1. *There is a full sub-SMC of $\mathscr{C}$, which we denote by $\mathscr{C}_K$, which is equivalent to* $\mathbb{R}^+$-Mat.
2. *The SMC $\mathscr{C}$ is enriched in commutative monoids*[7]*, and the enrichment on $\mathscr{C}_K$ coincides with the one given by the linear structure of* $\mathbb{R}^+$-Mat.
3. *The SMC $\mathscr{C}$ comes with an environment structure, i.e. with a family $(\dashv\!\Vert_{\mathscr{H}} : \mathscr{H} \to \mathbb{1})_{\mathscr{H}\in\mathrm{obj}\,\mathscr{C}}$ of morphisms which satisfy the following requirements:*

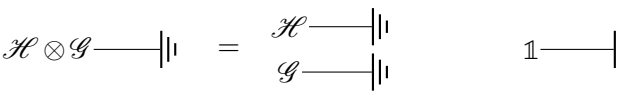

(3.1)

The environment structure on $\mathscr{C}_K$ coincides with the one given by the discarding maps of $\mathbb{R}^+$-Mat.

*The requirements above can be generalised by replacing $\mathbb{R}^+$ with an arbitrary commutative semiring R, in which case we will talk about an **R-probabilistic theory**.*

In the context of a specific R-probabilistic theory, we refer to $\mathscr{C}_K$ as the **classical theory**, to its object as the **classical systems** and to its morphisms as the **classical processes**. In particular, $\mathbb{1}$ is the tensor unit for $\mathscr{C}$, and the scalars of $\mathscr{C}$ form the semiring R. For each pair of objects $\mathscr{H},\mathscr{K} \in \mathrm{obj}\,\mathscr{C}$, the processes $\mathscr{H} \to \mathscr{K}$ in $\mathscr{C}$ form a commutative monoid, with a summation operation $+$ and a zero element 0 which we refer to as the **impossible process**. In line with the nomenclature adopted for classical systems, we refer to the $\dashv\!\Vert_{\mathscr{H}}$ maps involved in the environment structure as the **discarding maps**, and to those processes f satisfying the following equation as **normalised**:

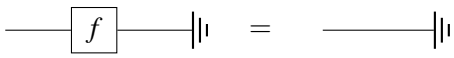

(3.2)

R-probabilistic theories come with a number of handy built-in features, which we now present.

[7] With linearity of both composition and tensor product.

3.2 Tests

Consider a process f with a classical output valued in some finite set Y:

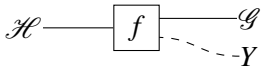

(3.3)

Operationally, we can think of each output value $y \in Y$ as corresponding to a different process being performed[8], and that process can be obtained by **testing against** the output value y:

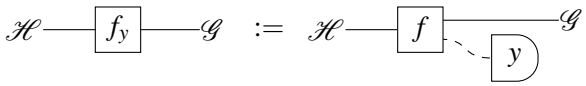

(3.4)

It should be noted that, even when f is normalised, the process f_y defined in Equation 3.4 is not necessarily so, instead being weighted by the probability of y occurring. In this sense, a process in the form of 3.3 corresponds to the notion of a **test** in OPTs [18,21]; if the input system $\mathscr{H}$ is trivial then it is a **preparation test**, while if the output system $\mathscr{K}$ is trivial then it is an **observation test**. Note that in this framework the difference between a test and a more general process is entirely in the eye of the beholder: when saying "test" instead of "process", we will mean that we are not interested in its classical inputs, but we will not exclude the presence of classical systems amongst the input systems of the process.

3.3 Output probability

The **probability**[9] $\mathbb{P}_\rho(y)$ of output value $y \in Y$ occurring in a preparation test can be obtained as follows:

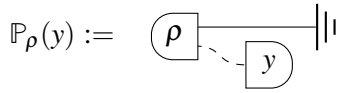

(3.5)

When the probability of outcome y for a preparation test ρ is invertible[10], we can **condition** on y to obtain the state $\frac{1}{\mathbb{P}_\rho(y)}\rho_y$, which is normalised whenever the original state ρ is:

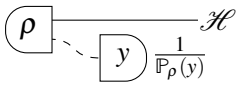

(3.6)

If ρ is normalised and the probabilities of all classical outcomes are either zero or invertible, then the reduced state $\rho|_{\mathscr{H}}$ can be written as usual in the form of a convex combination of the conditioned states:

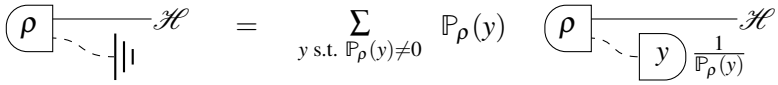

(3.7)

[8] A value y which is never output can, without loss of generality, be thought to correspond to the impossible process.

[9] We use the word loosely to intend the value in a generic commutative semiring R: as the latter is fixed by the choice of R-probabilistic theory, this should cause no confusion. In specific applications, one might want to use more specialised terms, such as **possibility** for $R = \mathbb{B}$, or **count** for $R = \mathbb{N}$.

[10] In the case of $\mathbb{R}^+$, this boils down to the usual requirement that the probability be non-vanishing.

3.4 Classical control

Processes in R-probabilistic theories are not limited to tests: in its most general form, a process f includes both some classical output set Y and some **classical input/control** set X:

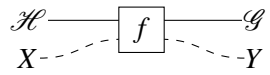

(3.8)

The process $f^{(x)}$ corresponding to a definite input value $x \in X$ is immediately obtained by applying f to the deterministic state corresponding to x:

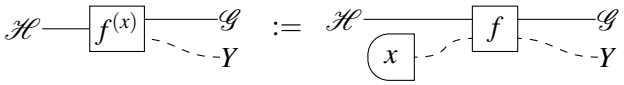

(3.9)

More in general, we can apply the process f above to an arbitrary state $p := (p_x)_{x \in X}$ on X to obtain a linear combination of the individual processes $f^{(x)}$ corresponding to each definite input value x:

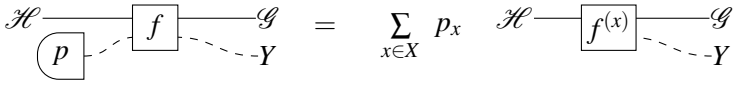

(3.10)

When the state is normalised, this gives a convex combination $\sum_x p_x f^{(x)}$ of the individual processes.

3.5 Preparations

As an example of the difference between the constructions above, we look at preparations. In an R-probabilistic theory, there are three possible notions of preparation: we have controlled preparations, preparation tests, and convex mixtures of the prepared states. A **controlled preparation** takes a classical input $x \in X$ and produces a state $f^{(x)}$ in $\mathscr{H}$, while a preparation test is a state on the joint system $\mathscr{H} \otimes X$:

(3.11)

controlled preparation preparation test

Given a normalised controlled preparation and a normalised state q on X, i.e. an R-distribution $(q_x)_{x \in X}$, we can always obtain a corresponding preparation test as follows:

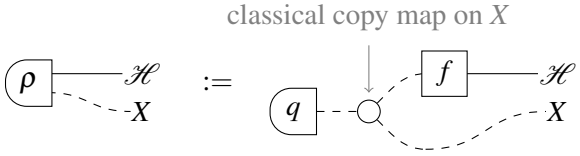

(3.12)

The normalisation requirement ensures that the probability of output $x \in X$ in the preparation test is the same specified by the probability distribution $(q_x)_{x \in X}$:

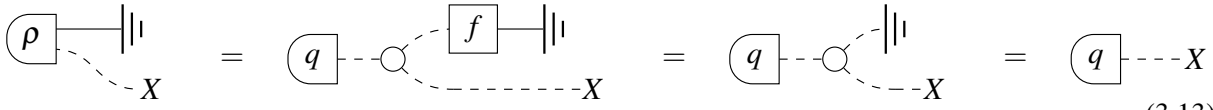

(3.13)

Given a preparation test obtained as in Equation 3.12, one can further discard the classical output system X to obtain the corresponding convex combination of prepared states, as shown in Equation 3.7. There is no compositional way of obtaining the controlled preparation back from the preparation test, or the preparation test back from the convex combination of prepared states.

3.6 Coarse-graining

Sometimes the classical outputs of a process f are not immediately interesting on their own, and one might wish to apply some form of **classical post-processing** to them, by which we mean the application of some additional classical (probabilistic) process to the output system. Most commonly, one obtains a new process g by applying a deterministic function $q : X \to Z$ to the relevant output system X of f, mapping the original outputs to some "more interesting" values in some other set Z:

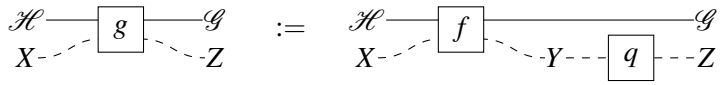

(3.14)

If we now focus on our new values in Z, we see that the process we obtain is exactly what the OPT framework refers to as a *coarse-graining* of process f:

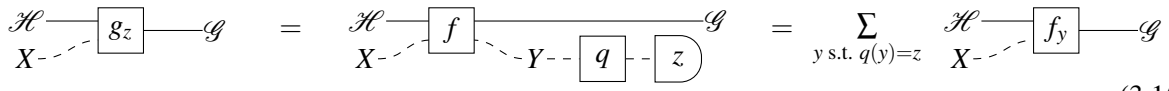

(3.15)

As a consequence, by a **coarse-graining** of a process f we will mean a process g taking the form of Equation 3.14 for some deterministic function $q : X \to Z$. Just like in the OPT framework, we say that f is a **refinement** of g whenever g is a coarse-graining of f.

3.7 Sharp preparations/observation pairs

A controlled preparation can be seen as a way of encoding the values of its classical input set X into its output system $\mathscr{H}$. However, not all preparations admit a corresponding decoding process which deterministically returns the original input values[11]: we refer to those preparations admitting one such decoding as **sharp preparations**, and they correspond to perfectly distinguishable states [18]. We say that a pair (p,m) of a controlled preparation $p : X \to \mathscr{H}$ and an observation test $m : \mathscr{H} \to X$ is a **sharp preparation/observation pair** (or **SPO pair**, for short) if the observation test m witnesses the sharpness of the controlled preparation p:

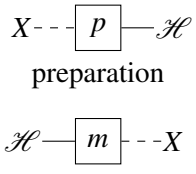

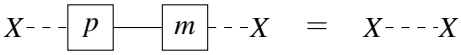

(3.16)

We say that an SPO pair is **normalised**[12] when both p and m are normalised processes. In fact, it suffices to ask for m to be normalised, as normalisation of p always follows[13]:

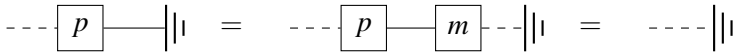

(3.17)

[11] The simplest example being given by non-injective deterministic functions.

[12] Normalised SPOs correspond to perfectly distinguishable states in OPTs.

[13] The converse is not in general true. As a simple counterexample, take a deterministic function $p : X \to Y$ which is injective but not surjective, and let $m : Y \rightharpoonup X$ be its left inverse, which is never total: then p is normalised, as it is a total function, but m is not total, and hence not normalised.

3.8 Decoherence maps

If (p,m) is a normalised SPO pair, we refer to the following process as the associated **decoherence map**:

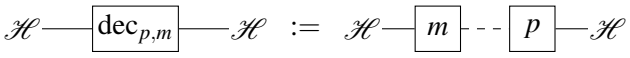

decoherence (3.18)

In order to keep track of the classical system X intervening in the decoherence map, we will use $K_{p,m}$ to denote it. There is good reason for our choice of name. Firstly, the process $\mathrm{dec}_{p,m}$ is normalised (because composition of normalised processes) and idempotent:

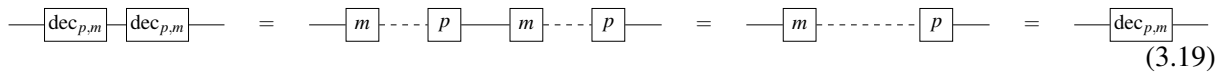

(3.19)

Secondly, decoherence maps can be used to reconstruct classical systems from arbitrary ones: to understand how this works, we use a mild variation on a common construction from category theory, known as *Karoubi envelope* (or *idempotent completion*). If $\mathscr{C}$ is a SMC with an environment structure, then we define the **normalised Karoubi envelope** of $\mathscr{C}$, denoted by $\mathrm{Split}_{\shortparallel}[\mathscr{C}]$, to be the SMC defined as follows:

- the systems are given by pairs in the form $(\mathscr{H},h)$, where $\mathscr{H}$ is a system of $\mathscr{C}$ and $h:\mathscr{H}\to\mathscr{H}$ is a normalised idempotent process;
- the processes $f:(\mathscr{H},h)\to(\mathscr{G},g)$ are exactly those processes $f:\mathscr{H}\to\mathscr{G}$ in $\mathscr{C}$ which are invariant under the specified idempotents, i.e. those which satisfy $f = g\circ f\circ h$;
- the tensor product on systems is given by $(\mathscr{H},h)\otimes(\mathscr{G},g) := (\mathscr{H}\otimes\mathscr{G}, h\otimes g)$, while the tensor product on processes is the one inherited from $\mathscr{C}$.

By looking at objects in the form $(\mathscr{H},id_{\mathscr{H}})$, it is immediate to see that $\mathscr{C}$ is a full sub-SMC of $\mathrm{Split}_{\shortparallel}[\mathscr{C}]$; in particular, the two categories have the same scalars. In fact, we can prove that $\mathrm{Split}_{\shortparallel}[\mathscr{C}]$ is an R-probabilistic theory whenever $\mathscr{C}$ is.

Lemma 2. *If $\mathscr{C}$ is an R-probabilistic theory, then so is* $\mathrm{Split}_{\shortparallel}[\mathscr{C}]$*, with classical systems, enrichment and discarding maps directly inherited from $\mathscr{C}$.*

Amongst the many objects of the normalised Karoubi envelope for an R-probabilistic theory sit all objects in the form $(\mathscr{H},\mathrm{dec}_{p,m})$, with (p,m) a normalised SPO pair on $\mathscr{H}$: we refer to systems in this form as **decohered systems**, and we now show that they can be used to recover the classical systems of the theory.

Theorem 3 (Decohered systems).
Let $\mathscr{C}$ be an R-probabilistic theory, and let $\mathrm{Decoh}[\mathscr{C}]$ *be the full subcategory of* $\mathrm{Split}_{\shortparallel}[\mathscr{C}]$ *spanned by the decohered systems. Then* $\mathrm{Decoh}[\mathscr{C}]$ *is a sub-SMC, and it is equivalent to the sub-SMC $\mathscr{C}_K$ of classical systems in $\mathscr{C}$. The equivalence sends a decohered system $(\mathscr{H},\mathrm{dec}_{p,m})$ to the classical system $K_{p,m}$, and a process $f:(\mathscr{H},\mathrm{dec}_{p,m})\to(\mathscr{G},\mathrm{dec}_{q,n})$ to the following classical process:*

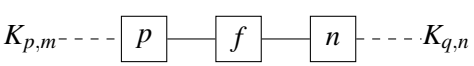

(3.20)

In the case of quantum theory, the decoherence maps associated with (possibly degenerate) observables arise from the normalised SPO pair given by considering the demolition measurement in the observable together with its adjoint. However, there are more decoherence maps in quantum theory than those associated with observables. Indeed, we have imposed no requirement for the observation to be adjoint to the preparation: hence decoherence maps are always normalised and idempotent, but not necessarily self-adjoint. As such, our notion of decoherence maps in quantum theory is more general than the self-adjoint normalised idempotent notion used in Ref. [52] and related works.

4 Quantum Theory

In CQM, dagger compact categories are taken to be abstract models of pure-state quantum theory, and the corresponding models of mixed-state quantum theory are obtained via the CPM construction [51]. The CPM category CPM[fdHilb] contains all positive states and completely positive maps of finite-dimensional quantum theory, it comes with an environment structure given by the trace, and it can be enriched in commutative monoids, with $\mathbb{R}^+$ as semiring of scalars (the usual $\mathbb{R}^+$-linear structure of CP maps). Unfortunately, CPM[fdHilb] does not come with a way of talking about classical systems, and in order to do so one usually appeals to the CP* construction [29, 35], which is related to C* algebras and quantum logic. In this work, we will not appeal to the CP* construction: instead, we will rely on the normalised Karoubi envelope to get a probabilistic theory out of CPM[fdHilb], following the steps of [52].

As we have seen before, the normalised Karoubi envelope $\mathrm{Split}_{\dashv\!\shortmid\!\shortmid}[\mathrm{CPM}[\mathrm{fdHilb}]]$ contains CPM[fdHilb] as a full sub-SMC, and we will refer to the objects in that sub-SMC as **quantum systems**. It also contains objects in the form $(\mathscr{H}, \mathrm{dec}_\bigcirc)$, where $\bigcirc$ is a special commutative †-Frobenius algebra on the finite-dimensional Hilbert space $\mathscr{H}$: we will refer to objects in this form as **classical systems**. The **decoherence map** $\mathrm{dec}_\bigcirc$ for $\bigcirc$ is the normalised idempotent process defined as follows:

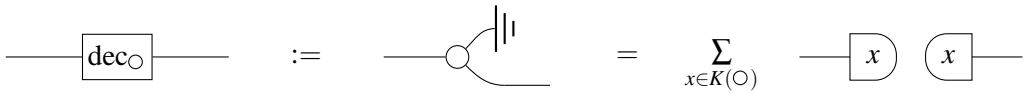

$$\sum_{x \in K(\bigcirc)} \tag{4.1}$$

where $(|x\rangle)_{x \in K(\bigcirc)}$ is the orthonormal basis formed by the classical states of $\bigcirc$ (see Ref. [34]).

Theorem 4 (**Classical systems** [52])**.**
Let $\mathscr{D}$ be the full subcategory of $\mathrm{Split}_{\dashv\!\shortmid\!\shortmid}[\mathrm{CPM}[\mathrm{fdHilb}]]$ *spanned by the classical systems. Then $\mathscr{D}$ is a full sub-SMC equivalent to* $\mathbb{R}^+$-Mat*. The equivalence sends a classical system $(\mathscr{H}, \mathrm{dec}_\bigcirc)$ to the set $K(\bigcirc)$ of classical states for $\bigcirc$, and a process $f : (\mathscr{H}, \mathrm{dec}_\bigcirc) \to (\mathscr{G}, \mathrm{dec}_\bullet)$ to the following classical process:*

$$\Big(\langle y|f|x\rangle\Big)_{(y,x) \in K(\bullet) \times K(\bigcirc)} : K(\bigcirc) \to K(\bullet) \tag{4.2}$$

When talking about **quantum theory** in the context of probabilistic theories, we will refer to the full sub-SMC of $\mathrm{Split}_{\dashv\!\shortmid\!\shortmid}[\mathrm{CPM}[\mathrm{fdHilb}]]$ jointly spanned by the quantum systems and the classical systems.

5 Causality and no-signalling

5.1 Causality

Causality is perhaps the trickiest point of contact between the OPT framework and probabilistic theories as we have defined them. On the one hand, causality is part of the very definition of probabilistic theories: a distinguished family of discarding maps is singled out by the basic operational requirement that states be localisable. This is how causality is usually understood in the CQM community [25, 32]. On the other hand, however, proving causality in a purely compositional way is impossible without reference to that very environment structure. Indeed, one encounters the following issue when trying to formulate the causality axiom as understood by the OPT community: the observation tests in OPTs have additional requirements that make them "normalised" in a suitable sense, while the ones defined in this work have no such requirement placed upon them. In order to single out a set of observation tests which would be suitable to formulate the causality axiom (as understood in OPTs), we would have to somehow refer to the environment structure to impose an appropriate normalisation condition[14].

[14] For example, note that CPM[fdHilb] admits environment structures different from the canonical one.

As a consequence of the remarks above, we will not talk about a *causality axiom* in the context of probabilistic theories. Instead, we will keep in mind the following general scenario, involving a (normalised) test f followed by a normalised process g:

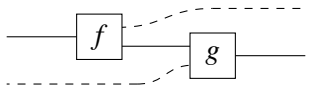

(5.1)

By **causality** we will mean the fact that the local output of the test f is independent of the input of the following process g, i.e. that normalised processes satisfy "no-signalling from the future" [18]:

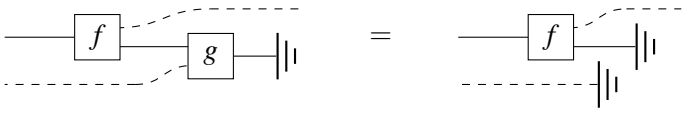

(5.2)

5.2 No-signalling

More in general, we define an N-party **Bell-type measurement scenario** [12, 47] in an R-probabilistic theory to be a process in the following form, where the processes $B_1, ..., B_N$ and the state ρ are all normalised:

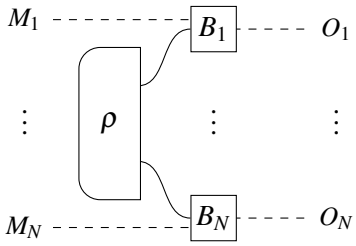

(5.3)

In the context of a Bell-type measurement scenario, the processes $B_1, ..., B_N$ are often referred to as **measurements**, their inputs as **measurement choices** and their outputs as **(measurement) outcomes**. We also refer to ρ as the **shared state**, to $\mathscr{M} := \prod_{j=1}^{N} M_j$ as the set of **joint measurement choices** and to $\prod_{j=1}^{N} O_j$ as the set of **joint (measurement) outcomes**.

Note that our definition is purely a matter of "shape", and makes no direct reference to probabilistic or no-signalling structure: one of the points of strength of our framework is indeed in providing that structure automatically, based on the shape of the process alone. Indeed, the Bell-type measurement scenarios defined above provide a direct point of contact with the sheaf-theoretic framework for non-locality and contextuality [1, 2] (in a way similar to, but much more direct than, what done in [23, 25]). In the sheaf-theoretic framework, a non-locality scenario is specified by a **no-signalling empirical model** $(\zeta_{\underline{m}})_{\underline{m} \in \mathscr{M}}$, which is a family of R-distributions $\zeta_{\underline{m}}$ on the set $\prod_{j=1}^{N} O_j$ of joint outcomes indexed by the joint measurement choices $\underline{m} \in \mathscr{M}$ (here R is the semiring chosen to model the desired notion of non-determinism).

Theorem 5 (**Bell-type measurement scenarios**)**.**
A Bell-type measurement scenario in an R-probabilistic theory always yields a no-signalling empirical model in the sheaf-theoretic framework, with R as semiring modelling non-determinism.

This result shows that the sheaf-theoretic framework can be applied directly to the treatment of non-locality in R-probabilistic theories, and conversely that R-probabilistic theories are a natural environment to study the realisability of no-signalling empirical models from the sheaf-theoretic framework.

6 Conclusions and future work

6.1 Conclusions

In this work, we have presented a general framework for probabilistic theories based upon three fundamental operational requirements: the existence of classical systems, the possibility of forming probabilistic combinations of processes (linear structure), and the possibility of discarding sub-systems (environment structure). Our requirements have a simple compositional and categorical formulation, as close as possible to the spirit and practice of Categorical Quantum Mechanics (CQM), and yet feature many of the fundamental structures used in Operational Probabilistic Theories (OPTs).

We hope that our work will contribute to providing a much-needed common ground for the CQM and OPT communities to collaborate on. In one direction, it brings the simplicity of CQM methods to OPTs, making it easy to create a new wealth of probabilistic theories. In the other direction, it allows methods and results from the OPT community to be straightforwardly applied to problems in CQM. Furthermore, we have shown that there is a direct connection between Bell tests in R-probabilistic theories and no-signalling empirical models in the sheaf-theoretic framework for non-locality and contextuality. As a consequence, methods from the sheaf-theoretic framework can be directly applied to problems of interest in CQM and OPTs, and conversely categorical methods can be used to study the realisability of empirical models in the sheaf-theoretic community.

6.2 Future work

Firstly, we have provided a new point of view on the emergence of classical systems within a probabilistic theory $\mathscr{C}$ (and within quantum theory in particular) based upon the *normalised* Karoubi envelope $\mathrm{Split}_{\dashv\Vdash}[\mathscr{C}]$. However, only a small subset of the objects of $\mathrm{Split}_{\dashv\Vdash}[\mathscr{C}]$ features in the construction. For example, it is known that all super-selected quantum systems appear in quantum theory by splitting all self-adjoint idempotent processes (i.e. using decoherence maps for both commutative and non-commutative special †-Frobenius algebras) [52], and this construction can be straightforwardly turned into a probabilistic theory. However, little or nothing is known about the operational significance of splitting general idempotent processes, and a thorough investigation will be the topic of future work.

Secondly, a number of theories independently developed in CQM and OPTs appear to be closely related, and should be properly connected. As a concrete example, we would be interest in applying our work to understand the connection between the ZW calculus [42] on the CQM side and Fermionic quantum theory [37, 38] on the OPT side. Similarly, work on thermodynamics has been undertaken by both communities [8, 14, 15, 16, 22], and should be related.

Thirdly, a future extended version of this work will contain a complete side-by-side translation of all common axioms from OPTs and constructions from CQM, for ease of use. The causality axiom has already been discussed, and two extremely important additional cases, namely the purity axioms from OPTs and the inner product structure from CQM, are covered in the Appendix.

Finally, some applications of our framework to the sheaf-theoretic study of non-locality and to quantum cryptography are soon to appear in Ref. [41]. In the future, we expect similar constructions to provide new insight into other families of empirical models of interest to the sheaf-theoretic and quantum cryptography communities.

Acknowledgements. The authors would like to thank Sean Tull, Matty Hoban and John Selby for comments, suggestions and useful discussions, as well as Sukrita Chatterji and Nicolò Chiappori for their support. SG gratefully acknowledges funding from EPSRC and the Williams Scholarship offered by Trinity College. CMS gratefully acknowledges funding from EPSRC and the Oxford-Google Deepmind Graduate Scholarship.

References

[1] Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal & Shane Mansfield (2015): *Contextuality, Cohomology and Paradox*. *arXiv preprint*, doi:10.4230/LIPIcs.CSL.2015.211. Available at `http://arxiv.org/abs/1502.03097`.

[2] Samson Abramsky & Adam Brandenburger (2011): *The sheaf-theoretic structure of non-locality and contextuality*. *New Journal of Physics* 13, doi:10.1088/1367-2630/13/11/113036.

[3] Samson Abramsky & Adam Brandenburger (2014): *An Operational Interpretation of Negative Probabilities and No-Signalling Models*. In: *Horizons of the Mind. A Tribute to Prakash Panangaden.*, pp. 59–75, doi:10.1007/978-3-319-06880-0_3. Available at `http://dx.doi.org/10.1007/978-3-319-06880-0{_}3http://link.springer.com/chapter/10.1007/978-3-319-06880-0{_}3`.

[4] Samson Abramsky & Bob Coecke (2009): *Categorical Quantum Mechanics*. *Handbook of Quantum Logic and Quantum Structures*, pp. 261–323, doi:10.1016/B978-0-444-52869-8.50010-4.

[5] Samson Abramsky & Lucien Hardy (2012): *Logical Bell inequalities*. *Phys. Rev. A* 85, p. 062114, doi:10.1103/PhysRevA.85.062114. Available at `http://link.aps.org/doi/10.1103/PhysRevA.85.062114`.

[6] Sabri W. Al-Safi & Anthony J. Short (2013): *Simulating all Nonsignaling Correlations via Classical or Quantum Theory with Negative Probabilities*. *Phys. Rev. Lett.* 111, p. 170403, doi:10.1103/PhysRevLett.111.170403. Available at `http://link.aps.org/doi/10.1103/PhysRevLett.111.170403`.

[7] Miriam Backens (2014): *The ZX-calculus is complete for stabilizer quantum mechanics*. *New Journal of Physics* 16(9), doi:10.1088/1367-2630/16/9/093021.

[8] Krzysztof Bar & Jamie Vicary (2014): *Groupoid Semantics for Thermal Computing*. *arXiv preprint*. Available at `http://arxiv.org/abs/1401.3280`.

[9] Howard Barnum, Jonathan Barrett, Matthew Leifer & Alexander Wilce (2007): *Generalized No-Broadcasting Theorem*. *Phys. Rev. Lett.* 99, p. 240501, doi:10.1103/PhysRevLett.99.240501.

[10] Howard Barnum & Alexander Wilce (2011): *Information Processing in Convex Operational Theories*. *Electronic Notes in Theoretical Computer Science* 270(1), pp. 3–15, doi:10.1016/j.entcs.2011.01.002. Proceedings of the Joint 5th International Workshop on Quantum Physics and Logic and 4th Workshop on Developments in Computational Models (QPL/DCM 2008).

[11] Jonathan Barrett (2007): *Information processing in generalized probabilistic theories*. *Physical Review A - Atomic, Molecular, and Optical Physics* 75(3), pp. 1–21, doi:10.1103/PhysRevA.75.032304.

[12] John S. Bell (1964): *On the Einstein-Podolsky-Rosen paradox*. *Physics* 1, pp. 195–200.

[13] Giulio Chiribella (2014): *Distinguishability and copiability of programs in general process theories*. *Int. J. Software Informatics* 8(3–4), pp. 209–223.

[14] Giulio Chiribella & Carlo Maria Scandolo (2015): *Operational axioms for diagonalizing states* 195, pp. 96–115. doi:10.4204/EPTCS.195.8.

[15] Giulio Chiribella & Carlo Maria Scandolo (2016): *Entanglement as an axiomatic foundation for statistical mechanics*. *arXiv:1608.04459 [quant-ph]*. Available at `http://arxiv.org/abs/1608.04459`.

[16] Giulio Chiribella & Carlo Maria Scandolo (2016): *Purity in microcanonical thermodynamics: a tale of three resource theories*. *arXiv:1608.04460 [quant-ph]*. Available at http://arxiv.org/abs/1608.04460.

[17] Giulio Chiribella (2014): *Dilation of states and processes in operational-probabilistic theories*. doi:10.4204/EPTCS.172.1. Available at http://arxiv.org/abs/1412.8539http://dx.doi.org/10.4204/EPTCS.172.1.

[18] Giulio Chiribella, Giacomo Mauro D'Ariano & Paolo Perinotti (2010): *Probabilistic theories with purification*. *Physical Review A - Atomic, Molecular, and Optical Physics* 81(6), doi:10.1103/PhysRevA.81.062348.

[19] Giulio Chiribella, Giacomo Mauro D'Ariano & Paolo Perinotti (2011): *Informational derivation of quantum theory*. *Physical Review A - Atomic, Molecular, and Optical Physics* 84(1), pp. 1–39, doi:10.1103/PhysRevA.84.012311.

[20] Giulio Chiribella, Giacomo Mauro D'Ariano & Paolo Perinotti (2012): *Quantum Theory, namely the pure and reversible theory of information*. *Entropy* 14(10), pp. 1877–1893, doi:10.3390/e14101877.

[21] Giulio Chiribella, Giacomo Mauro D'Ariano & Paolo Perinotti (2015): *Quantum from principles*. doi:10.1007/978-94-017-7303-4. Available at http://arxiv.org/abs/1506.00398http://dx.doi.org/10.1007/978-94-017-7303-4.

[22] Giulio Chiribella & Carlo Maria Scandolo (2015): *Entanglement and thermodynamics in general probabilistic theories*. *New Journal of Physics* 17(10), p. 103027, doi:10.1088/1367-2630/17/10/103027. Available at http://dx.doi.org/10.1088/1367-2630/17/10/103027.

[23] Giulio Chiribella & Xiao Yuan (2016): *Bridging the gap between general probabilistic theories and the device-independent framework for nonlocality and contextuality*. *Information and Computation*, doi:10.1016/j.ic.2016.02.006. Available at http://arxiv.org/abs/1504.02395{%}5Cnhttp://www.arxiv.org/pdf/1504.02395.pdf.

[24] Bob Coecke (2009): *Quantum Picturalism*. *Contemporary Physics*, pp. 1–32.

[25] Bob Coecke (2016): *Terminality Implies No-signalling ... and Much More Than That*. *New Generation Computing* 34, pp. 69–85.

[26] Bob Coecke & Ross Duncan (2011): *Interacting quantum observables: Categorical algebra and diagrammatics*. *New Journal of Physics* 13, doi:10.1088/1367-2630/13/4/043016.

[27] Bob Coecke & Bill Edwards (2012): *Spekkens's toy theory as a category of processes*. *Proceedings of Symposia in Applied Mathematics* 71, p. 28. Available at http://arxiv.org/abs/1108.1978{%}5Cnhttp://books.google.com/books?hl=en{&}lr={&}id=C3dllm0TU84C{&}oi=fnd{&}pg=PA61{&}dq=Spekkens's+toy+theory+as+a+category+of+processes{&}ots=1DIsn5vadm{&}sig=D7zK3kgNnMYvnreAoFEmI71Jypo.

[28] Bob Coecke & Chris Heunen (2012): *Pictures of complete positivity in arbitrary dimension*. *Electronic Proceedings in Theoretical Computer Science* 95, pp. 27–35, doi:10.4204/EPTCS.95.4.

[29] Bob Coecke, Chris Heunen & Aleks Kissinger (2014): *Categories of quantum and classical channels*. *Quantum Information Processing*, pp. 1–31, doi:10.1007/s11128-014-0837-4. Available at http://dx.doi.org/10.1007/s11128-014-0837-4.

[30] Bob Coecke & Aleks Kissinger (2015): *Categorical Quantum Mechanics I: Causal Quantum Processes*. *arXiv preprint*. Available at http://arxiv.org/abs/1510.05468.

[31] Bob Coecke & Aleks Kissinger (2016): *Picturing Quantum Processes*. Cambridge University Press.

[32] Bob Coecke & Raymond Lal (2013): *Causal Categories: Relativistically Interacting Processes*. *Foundations of Physics* 43(4), pp. 458–501, doi:10.1007/s10701-012-9646-8.

[33] Bob Coecke & Dusko Pavlovic (2008): *Quantum measurements without sums*. *arXiv preprint*. Available at http://arxiv.org/abs/quant-ph/0608035.

[34] Bob Coecke, Dusko Pavlovic & Jamie Vicary (2013): *A new description of orthogonal bases*. *Mathematical Structures in Computer Science* 23(03), doi:10.1017/S0960129512000047. Available at http://arxiv.org/abs/0810.0812.

[35] Oscar Cunningham & Chris Heunen (2015): *Axiomatizing complete positivity*. *Electronic Proceedings in Theoretical Computer Science* (Qpl 2015), pp. 148–157, doi:10.4204/EPTCS.195.11. Available at `http://arxiv.org/abs/1506.02931http://dx.doi.org/10.4204/EPTCS.195.11`.

[36] Borivoje Dakić & Časlav Brukner (2011): *Quantum theory and beyond: is entanglement special?* In Hans Halvorson, editor: *Deep Beauty: Understanding the Quantum World through Mathematical Innovation*, Cambridge University Press, Cambridge, pp. 365–392.

[37] Giacomo Mauro D'Ariano, Franco Manessi, Paolo Perinotti & Alessandro Tosini (2014): *Fermionic computation is non-local tomographic and violates monogamy of entanglement*. *EPL (Europhysics Letters)* 107(2), p. 20009, doi:10.1209/0295-5075/107/20009.

[38] Giacomo Mauro D'Ariano, Franco Manessi, Paolo Perinotti & Alessandro Tosini (2014): *The Feynman problem and fermionic entanglement: Fermionic theory versus qubit theory*. *International Journal of Modern Physics A* 29(17), p. 1430025, doi:10.1142/S0217751X14300257.

[39] Julien Degorre, Marc Kaplan, Sophie Laplante & Jérémie Roland (2009): *The Communication Complexity of Non-signaling Distributions*, pp. 270–281. Springer Berlin Heidelberg, Berlin, Heidelberg, doi:10.1007/978-3-642-03816-7_24. Available at `http://dx.doi.org/10.1007/978-3-642-03816-7_24`.

[40] Stefano Gogioso (2015): *A Bestiary of Sets and Relations*. *Electronic Proceedings in Theoretical Computer Science* (QPL 2015), pp. 208–227, doi:10.4204/EPTCS.195.16. Available at `http://arxiv.org/abs/1506.05025`.

[41] Stefano Gogioso & William Zeng (2017): *Generalised Mermin-type non-locality arguments*.

[42] Amar Hadzihasanovic (2015): *A diagrammatic axiomatisation for qubit entanglement*. *Proceedings - Symposium on Logic in Computer Science* 2015, pp. 573–584, doi:10.1109/LICS.2015.59.

[43] Lucien Hardy (2011): *Foliable operational structures for general probabilistic theories*. In H. Halvorson, editor: *Deep Beauty: Understanding the Quantum World through Mathematical Innovation*, Cambridge University Press, Cambridge, pp. 409–442.

[44] Lucien Hardy (2011): *Reformulating and reconstructing quantum theory*. *arXiv:1104.2066 [quant-ph]*. Available at `http://arxiv.org/abs/1104.2066`.

[45] Lucien Hardy (2016): *Quantum Theory: Informational Foundations and Foils*, chapter Reconstructing Quantum Theory, pp. 223–248. Springer Netherlands, Dordrecht, doi:10.1007/978-94-017-7303-4_7.

[46] Lucien Hardy (2001): *Quantum Theory From Five Reasonable Axioms*. *arXiv preprint*. Available at `http://arxiv.org/abs/quant-ph/0101012`.

[47] Joe Henson, Raymond Lal & Matthew F Pusey (2014): *Theory-independent limits on correlations from generalized Bayesian networks*. *New J. Phys.* 16(11), p. 113043, doi:10.1088/1367-2630/16/11/113043. Available at `http://stacks.iop.org/1367-2630/16/i=11/a=113043`.

[48] Clare Horsman (2011): *Quantum picturalism for topological cluster-state*. *New Journal of Physics* 133(9), doi:10.1088/1367-2630/13/9/095011.

[49] Aleks Kissinger (2012): *Pictures of Processes: Automated Graph Rewriting for Monoidal Categories and Applications to Quantum Computing*. Ph.D. thesis. Available at `http://arxiv.org/abs/1203.0202{%}5Cnhttp://kidkarolis.github.io/PolyChrome/{%}5Cnhttps://github.com/Quantomatic/quantomatic`.

[50] Lluis Masanes & Markus P. Müller (2011): *A derivation of quantum theory from physical requirements*. *New J. Phys.* 13(6), p. 063001. Available at `http://stacks.iop.org/1367-2630/13/i=6/a=063001`.

[51] Peter Selinger (2007): *Dagger Compact Closed Categories and Completely Positive Maps*. *Electronic Notes in Theoretical Computer Science* 170, pp. 139–163, doi:10.1016/j.entcs.2006.12.018.

[52] Peter Selinger (2008): *Idempotents in dagger categories (extended abstract)*. *Electronic Notes in Theoretical Computer Science* 210, pp. 107—-122, doi:10.1016/j.entcs.2008.04.021.

[53] Sean Tull (2016): *Operational Theories of Physics as Categories*. *arXiv preprint*. Available at `http://arxiv.org/abs/1602.06284`.

A Inner product structure

The category R-Mat always comes with compact closed structure, which takes the exact same form of the complex one from fdHilb $\simeq \mathbb{C}$-Mat. It can similarly be endowed with dagger structure, as long as the semiring R is involutive, i.e. it comes equipped with a self-inverse semiring homomorphism $\dagger : R \to R$. Every semiring can be turned into an involutive semiring by taking $\dagger := id_R$, and when talking about $\mathbb{R}^+$ in the context of a dagger structure we will give as understood that the involution is the identity.

Pure-state quantum theory notably comes with a non-degenerate inner product structure, which features heavily in the traditional formalisms. In the context of CQM, the inner product structure of pure-state quantum theory is captured by the notion of †-SMC, and intervenes in the construction of categories of completely positive maps [28, 51], the abstract counterparts to the operator model for mixed-state quantum theory. The non-degeneracy condition, on the other hand, can be captured directly in terms of the trace on positive operators, with no direct reference to the inner product structure itself: it says that the only positive operator ρ satisfying $\mathrm{Tr}\,\rho = 0$ is the zero operator, and more in general that the only CP map f satisfying $\mathrm{Tr}\,f = 0$ is the zero CP map. Because the trace provides the discarding maps for quantum theory, we will say that an R-probabilistic theory is **non-degenerate** if the following condition holds:

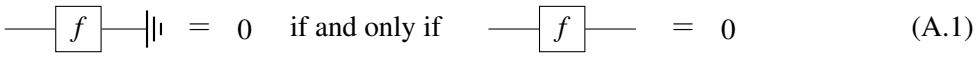

(A.1)

A necessary condition for non-degeneracy of an R-probabilistic theory is that the classical theory be itself non-degenerate, and the latter requirement is equivalent to positivity of the semiring R of scalars.

Lemma 6. *The classical theory* R-Mat *is non-degenerate if and only if* R *is a positive semiring.*

B Purity

By a pure process we mean one which cannot involve any non-trivial interaction with a discarded environment. To be precise, we will say that a process g is **pure** if the following implication always holds, for some normalised state ρ (which will, in general, depend on f) [13]:

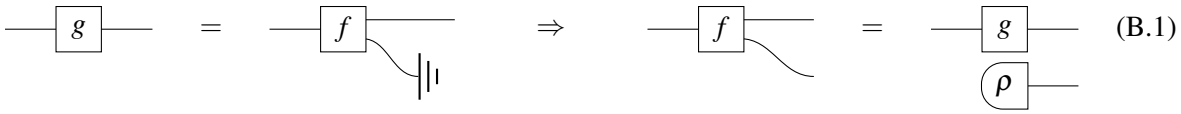

(B.1)

We will say that a R-probabilistic theory satisfies the **purification axiom** if every process g can be written in the following form for some pure process f (we refer to $\mathscr{E}$ as the **environment**) [17]:

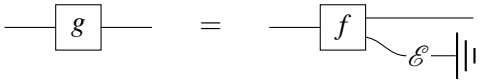

(B.2)

We will say that a probabilistic theory satisfies the **essentially unique purification axiom** [18, 20] if any two pure processes[15] $f, g : \mathscr{H} \to \mathscr{G} \otimes \mathscr{E}$ satisfy the following implication, for some normalised invertible process $u : \mathscr{E} \to \mathscr{E}$ (which will, in general, depend on both f and g):

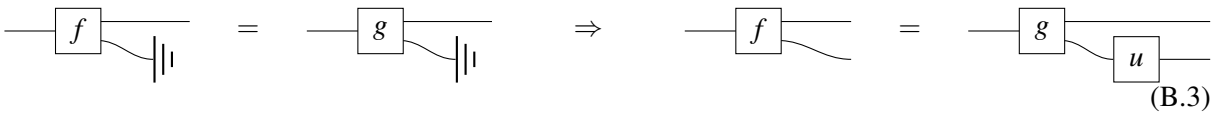

(B.3)

[15] Note that both processes have the same discarded environment system $\mathscr{E}$.

From a categorical perspective, it would be nice to phrase the *purity preservation* axiom from OPTs as the statement that pure processes form a sub-SMC, but unfortunately we encounter a difficulty with the explicit presence of classical systems in the theory. Indeed, the identities on non-trivial classical systems are not pure, as they can be obtained by discarding a branch of a classical copy operation. We could avoid this issue if no classical process were pure, but unfortunately this is not the case: for example, both the zero map and the partial functions with single-element domain are pure classical processes. However, it is true that there are no pure *normalised* classical processes aside from the classical deterministic states.

Lemma 7. *In R-*Mat*, the only pure normalised classical processes are the classical deterministic states.*

As a consequence, we will restrict our attention to pure normalised processes. We will say that a probabilistic theory satisfies the **normalised purity preservation axiom** if the pure normalised processes form a sub-SMC. Both classical theory and quantum theory satisfy the normalised purity preservation axiom.

C Proofs

Proof of Lemma 2 (p.9)

Because $\mathscr{C}$ is a full sub-SMC of Split$_{\dashv\Vdash}[\mathscr{C}]$, we already have classical systems available, and all we need to show is that the enrichment and environment structure can be suitably extended. The extension of the enrichment is taken care of by the observation that any linear combination of processes invariant under idempotents is itself invariant. Indeed, the distributivity of linear structure over composition in $\mathscr{C}$ yields:

$$\Big[\forall j.\ g\circ f_j\circ h = f_j\Big] \Rightarrow \Big[g\circ(\sum_j p_j f_j)\circ h = \sum_j p_j f_j\Big] \tag{C.1}$$

Furthermore, the same discarding maps from $\mathscr{C}$ yield an environment structure for Split$_{\dashv\Vdash}[\mathscr{C}]$: as the latter was defined in terms of *normalised* idempotents, the discarding maps themselves are always invariant. □

Proof of Theorem 3 (p.9)

All that we really need to show is that the process $f : (\mathscr{H}, \mathrm{dec}_{p,m}) \to (\mathscr{G}, \mathrm{dec}_{q,n})$ can be recovered from the classical process depicted in 3.20, and vice versa that any classical process $F : K_{p,m} \to K_{q,n}$ takes the form depicted in 3.20 for a unique process $f : (\mathscr{H}, \mathrm{dec}_{p,m}) \to (\mathscr{G}, \mathrm{dec}_{q,n})$. The following picture proves both directions:

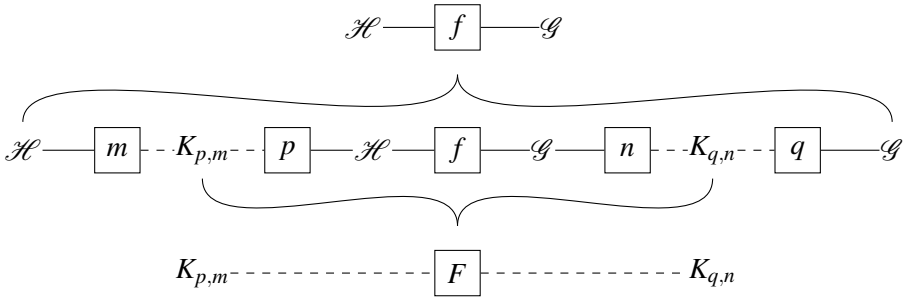

(C.2)

□

Proof of Theorem 5 (p.11)

We need to show two things:

(i) that each measurement choice $\underline{m} \in \prod_{j=1}^{N} M_j$ corresponds to an R-distribution on the joint measurement outcome set $\prod_{j=1}^{N} O_j$ (i.e. that the Bell-type measurement scenario corresponds to a well-defined conditions R-distribution);

(ii) that marginalising over the outcome set O_j for the j-th party yields an empirical model for the remaining parties which is independent of the individual measurement choice m_j made by the j-th party herself (i.e. that the conditional R-distribution satisfies the no-signalling condition).

We begin by showing that the no-signalling condition (ii) holds, and we obtain condition (i) as a corollary. Marginalising over the j-th outcome set in an R-probabilistic theory corresponds to hitting the corresponding classical system O_j with a discarding map:

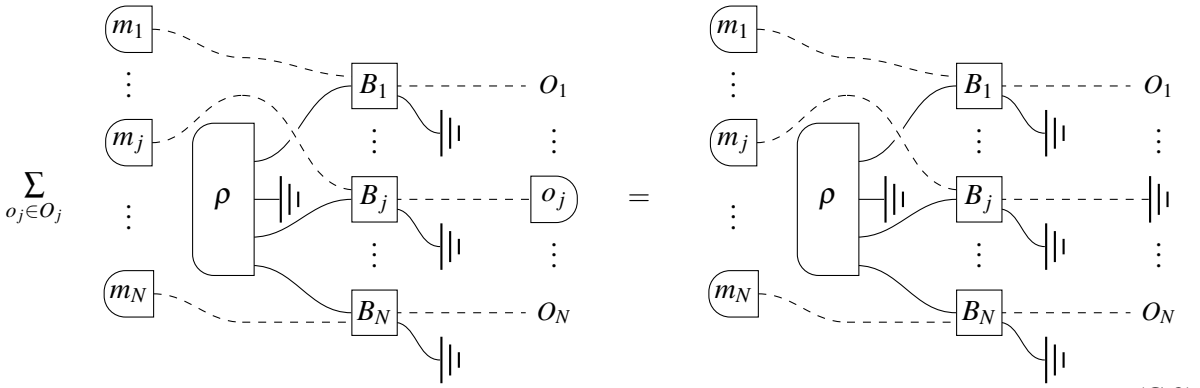

(C.3)

But the measurements in a Bell-type measurement scenario are required to be normalised, and hence this corresponds to an empirical model for the remaining parties which is independent of the choice of classical deterministic state $m_j \in M_j$ (which is always normalised, and hence discarded to the scalar 1):

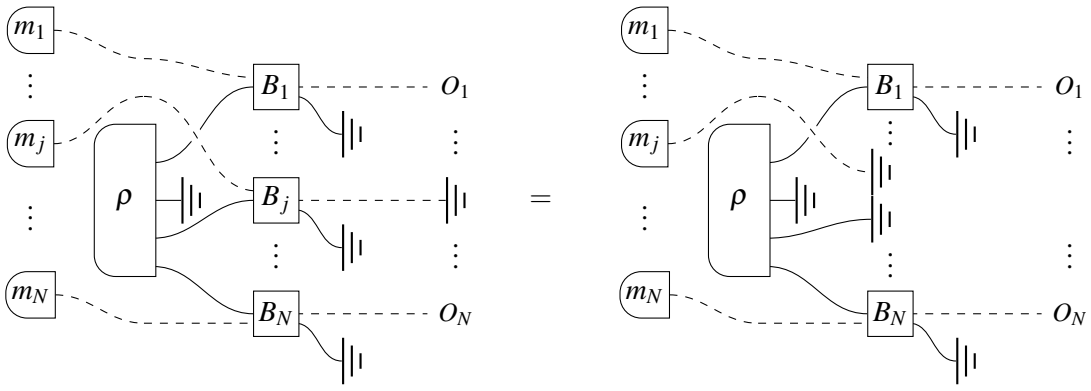

(C.4)

By repeating this process for all parties we are left with the scalar $\mathrm{Tr}\,\rho$: because the shared state ρ in a Bell-type measurement scenario is required to be normalised, this is the scalar 1. Hence the state over the joint measurement outcome set corresponding to the given joint measurement choice $\underline{m}$ was indeed normalised, i.e. we got an R-distribution as required by condition (i). □

Proof of Lemma 6 (p.16)

Because we can always evaluate classical processes on deterministic input states, it is sufficient to show that the following two conditions are equivalent: (i) the zero state is the only classical state which yields the zero scalar when hit with the discarding map; (ii) the semiring R of scalars is positive. A classical state on a classical system X takes the form $p := \sum_{x\in X} p_x \delta_x$, where δ_x is the deterministic classical state concentrated at $x \in X$, and $(p_x)_{x\in X}$ is some family of scalars: as a consequence, $\mathrm{Tr}\, p = \sum_{x\in X} p_x$ is a generic sum of scalars. The requirement of positivity for the semiring R states exactly that whenever $\sum_{x\in X} p_x = 0$ for some set X and some family $(p_x)_{x\in X}$ of scalars, then $p_x = 0$ for all $x \in X$: as a consequence, the requirement of positivity is equivalent to asking that the equation $\mathrm{Tr}\, p = 0$ holds exactly when the classical state p is the zero state $p := \sum_{x\in X} 0\delta_x$. □

Proof of Lemma 7 (p.17)

We can always classically copy the inputs/outputs of classical processes and then obtain the process back by discarding the copies: as a consequence, pure classical processes can have at most one input and one output. Furthermore, normalised classical processes must have the entire input system as domain. Hence normalised pure classical processes must have the tensor unit as input, and cover exactly one output (i.e. they cannot have non-trivial probability distributions as output, nor be the zero process): they are exactly the classical deterministic states. □